



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/808,185	03/24/2004	John C. Lalier	34826-1017	4569
41881 7590 03/17/2008 BRANDON N. SKLAR, ESQ. (PATENT PROSECUTION) KAYE SCHOLER, LLP 425 PARK AVENUE NEW YORK, NY 10022-3598				
EXAMINER DWTVEDI, MAHESH H				
ART UNIT		PAPER NUMBER		
2168				
MAIL DATE		DELIVERY MODE		
03/17/2008		PAPER		

**Please find below and/or attached an Office communication concerning this application or proceeding.**

The time period for reply, if any, is set in the attached communication.

### Office Action Summary

**Application No.**

10/808,185

**Applicant(s)**

LALLIER, JOHN C.

**Examiner**

MAHESH H. DWIVEDI

**Art Unit**

2168

**Period for Reply** -- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

**Status**

- 1) ☒ Responsive to communication(s) filed on 18 December 2007.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

**Disposition of Claims**

- 4) ☒ Claim(s) 19-26, 52-59, 79 and 88-90 is/are pending in the application.
- 4a) Of the above claim(s) \_\_\_\_\_ is/are withdrawn from consideration.
- 5) ☐ Claim(s) \_\_\_\_\_ is/are allowed.
- 6) ☒ Claim(s) 19-26, 52-59, 79 and 88-90 is/are rejected.
- 7) ☐ Claim(s) \_\_\_\_\_ is/are objected to.
- 8) ☐ Claim(s) \_\_\_\_\_ are subject to restriction and/or election requirement.

**Application Papers**

- 9) ☒ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 14 June 2004 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
- Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
- Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

**Priority under 35 U.S.C. § 119**

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some \* c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
  2. ☐ Certified copies of the priority documents have been received in Application No. \_\_\_\_\_.
  3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

\* See the attached detailed Office action for a list of the certified copies not received.

**Attachment(s)**

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☒ Information Disclosure Statement(s) (PTO/SB/08)  
Paper No(s)/Mail Date 11/08/2007.
- 4) ☐ Interview Summary (PTO-413)  
Paper No(s)/Mail Date \_\_\_\_\_.
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: \_\_\_\_\_.

## **DETAILED ACTION**

### ***Continued Examination Under 37 CFR 1.114***

1. A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 9/5/2006 has been entered.

### ***Election/Restrictions***

2. The examiner notes that in the reply to the restriction requirement filed on 12/18/2007, the applicant did not indicate whether the election was made with or without traverse. For the purposes of this examination, the examiner is assuming that the applicant elected without traverse.

3. Applicant's election without traverse of claims 19-26, 52-59, 79, and 88-90 in the reply filed on 12/18/2007 is acknowledged.

Claims 1-10, 11-18, 27-33, 34-43, 44-51, 60-66, 67-70, 72-75, 76-79, 81-87, and 90-91 are withdrawn from further consideration pursuant to 37 CFR 1.142(b) as being drawn to nonelected Invention I respectively, there being no allowable generic or linking claim. Election was made **without** traverse in the reply filed on 12/18/2007.

### ***Information Disclosure Statement***

4. The information disclosure statement (IDS) submitted on 11/08/2007 has been received, entered into the record, and considered. The submission is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

### ***Remarks***

5. Receipt of the Applicant's amendments filed on 08/30/2007 is acknowledged. The amendment includes the amending of claims 19, 24-25, 52, 54, 57-59, and 79, and the addition of claims 88-90.

### ***Specification***

Art Unit: 2168

6. The specification is objected to as failing to provide proper antecedent basis for the claimed subject matter. See 37 CFR 1.75(d)(1) and MPEP § 608.01(o). Correction of the following is required: Claim 89 recites the language of "pausing the migration procedure...resources" does not find any support in the specification.

***Claim Rejections - 35 USC § 112***

7. The following is a quotation of the first paragraph of 35 U.S.C. 112:

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

8. Claim 89 is rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the enablement requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to enable one skilled in the art to which it pertains, or with which it is most nearly connected, to make and/or use the invention. Specifically, the limitation of "pausing the migration procedure...resources" does not have any support from the specification.

***Claim Rejections - 35 USC § 102***

9. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(a) the invention was known or used by others in this country, or patented or described in a printed publication in this or a foreign country, before the invention thereof by the applicant for a patent.

10. Claims 19-21, 23-25, 52-54, 56-58, 79, 88, and 90 are rejected under 35 U.S.C. 102(a) as being anticipated by **Eshel et al.** (U.S. PG PUB 2003/0158862).

11. Regarding claim 19, **Eshel** teaches a method comprising:

A) storing in a target storage device a plurality of target data files corresponding respectively to respective ones of a plurality of source data files stored in a source storage device (Paragraph 130);

B) storing in each respective target data file information identifying the corresponding source data file identifying the corresponding source data file (Paragraph 132);

Art Unit: 2168

- C) activating a migration procedure to copy data from the source storage device to the target storage device, after target data files have been stored for all source data files in the plurality (Paragraph 130)
- D) receiving from a host device a request specifying a data file, while the migration procedure is executing (Paragraph 127);
- E) examining, in a target data file corresponding to the specified data file, selected information identifying a source data file (Paragraph 132);
- F) retrieving the requested data (Paragraph 127); and
- G) providing the requested data to the host device (Paragraph 127).

The examiner notes that **Eshel** teaches **“storing in a target storage device a plurality of target data files corresponding respectively to respective ones of a plurality of source data files stored in a source storage device”** as “These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system).” (Paragraph 130). The examiner further notes that **Eshel** teaches **“storing in each respective target data file information identifying the corresponding source data file identifying the corresponding source data file”** as “The exemplary embodiments of the present invention use snapshot tags to identify each snapshot and the file system from which that snapshot was captured. The snapshot tag notation used herein consists of the format (A:S1) wherein the first element, “A” in this example, identifies the file system and the second element, “S1” in this example, is the snapshot identifier for that snapshot. This allows the different file systems in the hot standby system described herein to capture snapshots at different times and only use a subset of those snapshots to synchronize the data between those file systems. The file systems of the exemplary embodiments generate a set of changes between snapshots that are captured for that file system. These sets of changes include a pair of tags to identify the snapshots between which the changes were determined. As an example, a snapshot tag pair (A:S2, A:S3) is included within a set of changes that were generated as the changes

that occurred between snapshot S2 and snapshot S3 that were captured on file system A. This set of changes is only able to be successfully applied to a file system that has been restored to the state of snapshot S2 from file system A. For example, if file system B receives this snapshot and snapshot S2 from file system A has not been restored to file system B or changes have not been applied to file system B that resulted in file system B having the state of snapshot (A:S2), application of the set of changes with the snapshot tag pair (A:S2,A:S3) is inappropriate. A file system discards a set of changes that is received and does not have a snapshot pair that starts with a snapshot tag that corresponds to the most recently restored or applied snapshot to that file system. Exemplary systems identify the last applied or restored snapshot and request from the other file system the set of changes that corresponds to the changes made since the last applied or restored snapshot" (Paragraph 132). The examiner further notes that Eshel teaches "activating a migration procedure to copy data from the source storage device to the target storage device, after target data files have been stored for all source data files in the plurality" as These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system). These embodiments then periodically bring the standby or mirror file system up-to-date by generating new snapshots of the original file system and determining the changes between these new, more recently captured or generated snapshots and the state that was captured by a previous snapshot of the original file system that had been transferred to the mirror file system. The original file system generates a set of changes that are then communicated and applied to the standby file system in order to bring the standby file system up to the state of the new snapshots captured on the original file system" (Paragraph 130). The examiner further notes that Eshel teaches "receiving from a host device a request specifying a data file, while the migration procedure is executing" as "Another common use of snapshots is to back up a file system to tape while allowing continued read/write access to the file system during the backup process" (Paragraph 127). The examiner further notes that Eshel teaches "examining, in a

**target data file corresponding to the specified data file, selected information**

**identifying a source data file** as “The exemplary embodiments of the present invention use snapshot tags to identify each snapshot and the file system from which that snapshot was captured. The snapshot tag notation used herein consists of the format (A:S1) wherein the first element, “A” in this example, identifies the file system and the second element, “S1” in this example, is the snapshot identifier for that snapshot. This allows the different file systems in the hot standby system described herein to capture snapshots at different times and only use a subset of those snapshots to synchronize the data between those file systems. The file systems of the exemplary embodiments generate a set of changes between snapshots that are captured for that file system. These sets of changes include a pair of tags to identify the snapshots between which the changes were determined. As an example, a snapshot tag pair (A:S2, A:S3) is included within a set of changes that were generated as the changes that occurred between snapshot S2 and snapshot S3 that were captured on file system A. This set of changes is only able to be successfully applied to a file system that has been restored to the state of snapshot S2 from file system A. For example, if file system B receives this snapshot and snapshot S2 from file system A has not been restored to file system B or changes have not been applied to file system B that resulted in file system B having the state of snapshot (A:S2), application of the set of changes with the snapshot tag pair (A:S2,A:S3) is inappropriate. A file system discards a set of changes that is received and does not have a snapshot pair that starts with a snapshot tag that corresponds to the most recently restored or applied snapshot to that file system. Exemplary systems identify the last applied or restored snapshot and request from the other file system the set of changes that corresponds to the changes made since the last applied or restored snapshot” (Paragraph 132). The examiner further notes that Eshel teaches “**retrieving the requested data**” as “Another common use of snapshots is to back up a file system to tape while allowing continued read/write access to the file system during the backup process” (Paragraph 127). The examiner further notes that Eshel teaches “**providing the requested data to the host device**” as “Another

common use of snapshots is to back up a file system to tape while allowing continued read/write access to the file system during the backup process" (Paragraph 127).

Regarding claim 20, **Eshel** further teaches a method comprising:

A) wherein the source data file is stored in a file volume on the source storage device (Paragraph 129, Figure 15a).

The examiner notes that **Eshel** teaches "**wherein the source data file is stored in a file volume on the source storage device**" as "A block diagram of an overall system architecture for a primary and standby file system 1500 according to an exemplary embodiment of the present invention is illustrated in FIG. 15A. This exemplary system architecture has a primary file system, denoted as file system A 1502, a standby file system, denoted as file system B 1504 and a network 106 to provide communications between these file systems. Alternative embodiments maintain the primary and backup file systems within a single processor, thereby obviating the requirement for a network 106. File system A 1502 in this example has two snapshot datasets, a first snapshot dataset 1506 and a second snapshot dataset 1508. These two snapshot datasets captured the state of the file system A 1502 at different times. File system A 1502 operates by communicating snapshot datasets, such as first snapshot dataset 1506 and second snapshot 1508, to file system B 1504. File system B 1504, in turn, stores copies of the snapshot datasets that are received from file system A 1502. File system B 1504 stores a first snapshot dataset copy 1510 and a second snapshot dataset copy 1512 to support standby data storage operations" (Paragraph 129).

Regarding claim 21, **Eshel** further teaches a method comprising:

A) wherein the target data file is stored in a file volume on the target storage device (Paragraph 129, Figure 15a).

The examiner notes that **Eshel** teaches "**wherein the target data file is stored in a file volume on the target storage device**" as "A block diagram of an overall system architecture for a primary and standby file system 1500 according to an



exemplary embodiment of the present invention is illustrated in FIG. 15A. This exemplary system architecture has a primary file system, denoted as file system A 1502, a standby file system, denoted as file system B 1504 and a network 106 to provide communications between these file systems. Alternative embodiments maintain the primary and backup file systems within a single processor, thereby obviating the requirement for a network 106. File system A 1502 in this example has two snapshot datasets, a first snapshot dataset 1506 and a second snapshot dataset 1508. These two snapshot datasets captured the state of the file system A 1502 at different times. File system A 1502 operates by communicating snapshot datasets, such as first snapshot dataset 1506 and second snapshot 1508, to file system B 1504. File system B 1504, in turn, stores copies of the snapshot datasets that are received from file system A 1502. File system B 1504 stores a first snapshot dataset copy 1510 and a second snapshot dataset copy 1512 to support standby data storage operations" (Paragraph 129).

Regarding claim 23, **Eshel** further teaches a method comprising:

A) wherein the target storage device comprises a file server (Paragraph 49).

The examiner notes that **Eshel** teaches "**wherein the target data file is stored in a file volume on the target storage device**" as "In another embodiment of the present invention, the computer systems of file system 102 and backup 108 are a server such as one or more computers executing operating systems such as SunOS or AIX, such as SUN Ultra workstations running the SunOS operating system or IBM RS/6000 workstations and servers running the AIX operating system" (Paragraph 49).

Regarding claim 24, **Eshel** further teaches a method comprising:

A) wherein the request is received from the host device via a network (Paragraph 129).

The examiner notes that **Eshel** teaches "**wherein the request is received from the host device via a network**" as "A block diagram of an overall system architecture for a primary and standby file system 1500 according to an exemplary embodiment of the present invention is illustrated in FIG. 15A. This exemplary system architecture has

a primary file system, denoted as file system A 1502, a standby file system, denoted as file system B 1504 and a network 106 to provide communications between these file systems" (Paragraph 129).

Regarding claim 25, **Eshel** further teaches a method comprising:

A) wherein the selected information in a respective target data file identifies a logical location of the corresponding source data file in a source file volume (Paragraph 132).

The examiner notes that **Eshel** teaches "**wherein the selected information in a respective target data file identifies a logical location of the corresponding source data file in a source file volume**" as "The exemplary embodiments of the present invention use snapshot tags to identify each snapshot and the file system from which that snapshot was captured. The snapshot tag notation used herein consists of the format (A:S1) wherein the first element, "A" in this example, identifies the file system and the second element, "S1" in this example, is the snapshot identifier for that snapshot. This allows the different file systems in the hot standby system described herein to capture snapshots at different times and only use a subset of those snapshots to synchronize the data between those file systems. The file systems of the exemplary embodiments generate a set of changes between snapshots that are captured for that file system. These sets of changes include a pair of tags to identify the snapshots between which the changes were determined. As an example, a snapshot tag pair (A:S2, A:S3) is included within a set of changes that were generated as the changes that occurred between snapshot S2 and snapshot S3 that were captured on file system A. This set of changes is only able to be successfully applied to a file system that has been restored to the state of snapshot S2 from file system A. For example, if file system B receives this snapshot and snapshot S2 from file system A has not been restored to file system B or changes have not been applied to file system B that resulted in file system B having the state of snapshot (A:S2), application of the set of changes with the snapshot tag pair (A:S2,A:S3) is inappropriate. A file system discards a set of changes that is received and does not have a snapshot pair that starts with a snapshot tag that corresponds to the most recently restored or applied snapshot to that file system.

Exemplary systems identify the last applied or restored snapshot and request from the other file system the set of changes that corresponds to the changes made since the last applied or restored snapshot" (Paragraph 132).

Regarding claim 52, **Eshel** teaches a system comprising:

- A) a target storage device configured to store data files (Paragraph 130);
- B) a source storage device configured to store data files (Paragraph 130);
- C) at least one processor configured to: store in the target storage device a plurality of target data files corresponding respectively to respective ones of a plurality of source data files in the source storage device (Paragraph 130);
- D) store in each respective target data file information identifying the corresponding source data file (Paragraph 132);
- E) activate a migration procedure to copy data from the source storage device to the target storage device, after the target data files have been stored for all source data files in the plurality (Paragraph 130); and
- F) an interface configured to: receive from a host device a request specifying a data file, while the migration procedure is executing (Paragraph 127);
- G) wherein the least one processor is further configured to: examiner, a target data file corresponding to the specified data file, selected information identifying a source data file (Paragraph 132);
- H) retrieve requested data (Paragraph 127); and
- I) provide the requested data to the host device (Paragraph 127).

The examiner notes that **Eshel** teaches "a target storage device configured to store data files" as "These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system)." The examiner further notes that **Eshel** teaches "a source storage device configured to store data files" as "These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and

transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system)." The examiner further notes that Eshel teaches **"at least one processor configured to: store in the target storage device a plurality of target data files corresponding respectively to respective ones of a plurality of source data files in the source storage device"** as "These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system)." The examiner further notes that Eshel teaches **"store in each respective target data file information identifying the corresponding source data file"** as "The exemplary embodiments of the present invention use snapshot tags to identify each snapshot and the file system from which that snapshot was captured. The snapshot tag notation used herein consists of the format (A:S1) wherein the first element, "A" in this example, identifies the file system and the second element, "S1" in this example, is the snapshot identifier for that snapshot. This allows the different file systems in the hot standby system described herein to capture snapshots at different times and only use a subset of those snapshots to synchronize the data between those file systems. The file systems of the exemplary embodiments generate a set of changes between snapshots that are captured for that file system. These sets of changes include a pair of tags to identify the snapshots between which the changes were determined. As an example, a snapshot tag pair (A:S2, A:S3) is included within a set of changes that were generated as the changes that occurred between snapshot S2 and snapshot S3 that were captured on file system A. This set of changes is only able to be successfully applied to a file system that has been restored to the state of snapshot S2 from file system A. For example, if file system B receives this snapshot and snapshot S2 from file system A has not been restored to file system B or changes have not been applied to file system B that resulted in file system B having the state of snapshot (A:S2), application of the set of changes with the snapshot tag pair (A:S2,A:S3) is inappropriate. A file system discards a set of changes that is received and does not have a snapshot pair that starts

with a snapshot tag that corresponds to the most recently restored or applied snapshot to that file system. Exemplary systems identify the last applied or restored snapshot and request from the other file system the set of changes that corresponds to the changes made since the last applied or restored snapshot" (Paragraph 132). The examiner notes that Eshel teaches **"activate a migration procedure to copy data from the source storage device to the target storage device, after the target data files have been stored for all source data files in the plurality"** as "These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system). These embodiments then periodically bring the standby or mirror file system up-to-date by generating new snapshots of the original file system and determining the changes between these new, more recently captured or generated snapshots and the state that was captured by a previous snapshot of the original file system that had been transferred to the mirror file system. The original file system generates a set of changes that are then communicated and applied to the standby file system in order to bring the standby file system up to the state of the new snapshots captured on the original file system" (Paragraph 130). The examiner further notes that Eshel teaches **"an interface configured to: receive from a host device a request specifying a data file, while the migration procedure is executing"** as "Another common use of snapshots is to back up a file system to tape while allowing continued read/write access to the file system during the backup process" (Paragraph 127). The examiner further notes that Eshel teaches **"wherein the least one processor is further configured to: examiner, a target data file corresponding to the specified data file, selected information identifying a source data file"** as "The exemplary embodiments of the present invention use snapshot tags to identify each snapshot and the file system from which that snapshot was captured. The snapshot tag notation used herein consists of the format (A:S1) wherein the first element, "A" in this example, identifies the file system and the second element, "S1" in this example, is the snapshot identifier for that snapshot. This allows the different file systems in the hot standby

system described herein to capture snapshots at different times and only use a subset of those snapshots to synchronize the data between those file systems. The file systems of the exemplary embodiments generate a set of changes between snapshots that are captured for that file system. These sets of changes include a pair of tags to identify the snapshots between which the changes were determined. As an example, a snapshot tag pair (A:S2, A:S3) is included within a set of changes that were generated as the changes that occurred between snapshot S2 and snapshot S3 that were captured on file system A. This set of changes is only able to be successfully applied to a file system that has been restored to the state of snapshot S2 from file system A. For example, if file system B receives this snapshot and snapshot S2 from file system A has not been restored to file system B or changes have not been applied to file system B that resulted in file system B having the state of snapshot (A:S2), application of the set of changes with the snapshot tag pair (A:S2,A:S3) is inappropriate. A file system discards a set of changes that is received and does not have a snapshot pair that starts with a snapshot tag that corresponds to the most recently restored or applied snapshot to that file system. Exemplary systems identify the last applied or restored snapshot and request from the other file system the set of changes that corresponds to the changes made since the last applied or restored snapshot" (Paragraph 132). The examiner further notes that **Eshel** teaches "**retrieving the requested data**" as "Another common use of snapshots is to back up a file system to tape while allowing continued read/write access to the file system during the backup process" (Paragraph 127). The examiner further notes that **Eshel** teaches "**providing the requested data to the host device**" as "Another common use of snapshots is to back up a file system to tape while allowing continued read/write access to the file system during the backup process" (Paragraph 127).

Regarding claim 53, **Eshel** further teaches a system comprising:

A) wherein the source data file is stored in a file volume on the source storage device (Paragraph 129, Figure 15a).

The examiner notes that **Eshel** teaches “**wherein the source data file is stored in a file volume on the source storage device**” as “A block diagram of an overall system architecture for a primary and standby file system 1500 according to an exemplary embodiment of the present invention is illustrated in FIG. 15A. This exemplary system architecture has a primary file system, denoted as file system A 1502, a standby file system, denoted as file system B 1504 and a network 106 to provide communications between these file systems. Alternative embodiments maintain the primary and backup file systems within a single processor, thereby obviating the requirement for a network 106. File system A 1502 in this example has two snapshot datasets, a first snapshot dataset 1506 and a second snapshot dataset 1508. These two snapshot datasets captured the state of the file system A 1502 at different times. File system A 1502 operates by communicating snapshot datasets, such as first snapshot dataset 1506 and second snapshot 1508, to file system B 1504. File system B 1504, in turn, stores copies of the snapshot datasets that are received from file system A 1502. File system B 1504 stores a first snapshot dataset copy 1510 and a second snapshot dataset copy 1512 to support standby data storage operations” (Paragraph 129).

Regarding claim 54, **Eshel** further teaches a system comprising:

A) wherein the target data file is stored in a file volume on the target storage device (Paragraph 129, Figure 15a).

The examiner notes that **Eshel** teaches “**wherein the target data file is stored in a file volume on the target storage device**” as “A block diagram of an overall system architecture for a primary and standby file system 1500 according to an exemplary embodiment of the present invention is illustrated in FIG. 15A. This exemplary system architecture has a primary file system, denoted as file system A 1502, a standby file system, denoted as file system B 1504 and a network 106 to provide communications between these file systems. Alternative embodiments maintain the primary and backup file systems within a single processor, thereby obviating the requirement for a network 106. File system A 1502 in this example has two snapshot

datasets, a first snapshot dataset 1506 and a second snapshot dataset 1508. These two snapshot datasets captured the state of the file system A 1502 at different times. File system A 1502 operates by communicating snapshot datasets, such as first snapshot dataset 1506 and second snapshot 1508, to file system B 1504. File system B 1504, in turn, stores copies of the snapshot datasets that are received from file system A 1502. File system B 1504 stores a first snapshot dataset copy 1510 and a second snapshot dataset copy 1512 to support standby data storage operations” (Paragraph 129).

Regarding claim 56, **Eshel** further teaches a system comprising:

A) wherein the target storage device comprises a file server (Paragraph 49).

The examiner notes that **Eshel** teaches “**wherein the target data file is stored in a file volume on the target storage device**” as “In another embodiment of the present invention, the computer systems of file system 102 and backup 108 are a server such as one or more computers executing operating systems such as SunOS or AIX, such as SUN Ultra workstations running the SunOS operating system or IBM RS/6000 workstations and servers running the AIX operating system” (Paragraph 49).

Regarding claim 57, **Eshel** further teaches a system comprising:

A) wherein the request is received from the host device via a network (Paragraph 129).

The examiner notes that **Eshel** teaches “**wherein the request is received from the host device via a network**” as “A block diagram of an overall system architecture for a primary and standby file system 1500 according to an exemplary embodiment of the present invention is illustrated in FIG. 15A. This exemplary system architecture has a primary file system, denoted as file system A 1502, a standby file system, denoted as file system B 1504 and a network 106 to provide communications between these file systems” (Paragraph 129).

Regarding claim 58, **Eshel** further teaches a system comprising:



Art Unit: 2168

A) wherein the selected information in a respective target data file identifies a logical location of the corresponding source data file in a source file volume (Paragraph 132).

The examiner notes that **Eshel** teaches “**wherein the selected information in a respective target data file identifies a logical location of the corresponding source data file in a source file volume**” as “The exemplary embodiments of the present invention use snapshot tags to identify each snapshot and the file system from which that snapshot was captured. The snapshot tag notation used herein consists of the format (A:S1) wherein the first element, “A” in this example, identifies the file system and the second element, “S1” in this example, is the snapshot identifier for that snapshot. This allows the different file systems in the hot standby system described herein to capture snapshots at different times and only use a subset of those snapshots to synchronize the data between those file systems. The file systems of the exemplary embodiments generate a set of changes between snapshots that are captured for that file system. These sets of changes include a pair of tags to identify the snapshots between which the changes were determined. As an example, a snapshot tag pair (A:S2, A:S3) is included within a set of changes that were generated as the changes that occurred between snapshot S2 and snapshot S3 that were captured on file system A. This set of changes is only able to be successfully applied to a file system that has been restored to the state of snapshot S2 from file system A. For example, if file system B receives this snapshot and snapshot S2 from file system A has not been restored to file system B or changes have not been applied to file system B that resulted in file system B having the state of snapshot (A:S2), application of the set of changes with the snapshot tag pair (A:S2,A:S3) is inappropriate. A file system discards a set of changes that is received and does not have a snapshot pair that starts with a snapshot tag that corresponds to the most recently restored or applied snapshot to that file system. Exemplary systems identify the last applied or restored snapshot and request from the other file system the set of changes that corresponds to the changes made since the last applied or restored snapshot” (Paragraph 132).

Regarding claim 79, **Eshel** further teaches a method comprising:

Art Unit: 2168

A) copying the identified source data file from the source storage device to the target storage device (Paragraph 49).

The examiner notes that **Eshel** teaches “**copying the identified source data file from the source storage device to the target storage device**” as “These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system). These embodiments then periodically bring the standby or mirror file system up-to-date by generating new snapshots of the original file system and determining the changes between these new, more recently captured or generated snapshots and the state that was captured by a previous snapshot of the original file system that had been transferred to the mirror file system. The original file system generates a set of changes that are then communicated and applied to the standby file system in order to bring the standby file system up to the state of the new snapshots captured on the original file system” (Paragraph 130).

Regarding claim 88, **Eshel** further teaches a method comprising:

A) activating a migration procedure to copy source data files from the source storage device to locations of corresponding target data files (Paragraph 49).

The examiner notes that **Eshel** teaches “**activating a migration procedure to copy source data files from the source storage device to locations of corresponding target data files**” as “These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system). These embodiments then periodically bring the standby or mirror file system up-to-date by generating new snapshots of the original file system and determining the changes between these new, more recently captured or generated snapshots and the state that was captured by a previous snapshot of the original file system that had been transferred to the mirror file system. The original file system generates a set of changes

Art Unit: 2168

that are then communicated and applied to the standby file system in order to bring the standby file system up to the state of the new snapshots captured on the original file system" (Paragraph 130).

Regarding claim 90, **Eshel** teaches a method comprising:

- A) storing in a target storage device a plurality of target data files corresponding respectively to respective ones of a plurality of source data files stored in a source storage device (Paragraphs 130 and 132);
- B) storing in each respective target data file information identifying the corresponding source data file (Paragraph 132); and
- C) activating a migration procedure to copy source data files from the source storage device to locations of the corresponding target data files in the target storage device (Paragraph 130).

The examiner notes that **Eshel** teaches **"storing in a target storage device a plurality of target data files corresponding respectively to respective ones of a plurality of source data files stored in a source storage device"** as "These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system)." (Paragraph 130) and "The exemplary embodiments of the present invention use snapshot tags to identify each snapshot and the file system from which that snapshot was captured. The snapshot tag notation used herein consists of the format (A:S1) wherein the first element, "A" in this example, identifies the file system and the second element, "S1" in this example, is the snapshot identifier for that snapshot. This allows the different file systems in the hot standby system described herein to capture snapshots at different times and only use a subset of those snapshots to synchronize the data between those file systems. The file systems of the exemplary embodiments generate a set of changes between snapshots that are captured for that file system. These sets of changes include a pair of tags to identify the snapshots between which the changes were determined. As an example, a

snapshot tag pair (A:S2, A:S3) is included within a set of changes that were generated as the changes that occurred between snapshot S2 and snapshot S3 that were captured on file system A. This set of changes is only able to be successfully applied to a file system that has been restored to the state of snapshot S2 from file system A. For example, if file system B receives this snapshot and snapshot S2 from file system A has not been restored to file system B or changes have not been applied to file system B that resulted in file system B having the state of snapshot (A:S2), application of the set of changes with the snapshot tag pair (A:S2,A:S3) is inappropriate. A file system discards a set of changes that is received and does not have a snapshot pair that starts with a snapshot tag that corresponds to the most recently restored or applied snapshot to that file system. Exemplary systems identify the last applied or restored snapshot and request from the other file system the set of changes that corresponds to the changes made since the last applied or restored snapshot" (Paragraph 132). The examiner further notes that Eshel teaches **"storing in each respective target data file information identifying the corresponding source data file"** as "The exemplary embodiments of the present invention use snapshot tags to identify each snapshot and the file system from which that snapshot was captured. The snapshot tag notation used herein consists of the format (A:S1) wherein the first element, "A" in this example, identifies the file system and the second element, "S1" in this example, is the snapshot identifier for that snapshot. This allows the different file systems in the hot standby system described herein to capture snapshots at different times and only use a subset of those snapshots to synchronize the data between those file systems. The file systems of the exemplary embodiments generate a set of changes between snapshots that are captured for that file system. These sets of changes include a pair of tags to identify the snapshots between which the changes were determined. As an example, a snapshot tag pair (A:S2, A:S3) is included within a set of changes that were generated as the changes that occurred between snapshot S2 and snapshot S3 that were captured on file system A. This set of changes is only able to be successfully applied to a file system that has been restored to the state of snapshot S2 from file system A. For example, if file system B receives this snapshot and snapshot S2 from file system A has

not been restored to file system B or changes have not been applied to file system B that resulted in file system B having the state of snapshot (A:S2), application of the set of changes with the snapshot tag pair (A:S2,A:S3) is inappropriate. A file system discards a set of changes that is received and does not have a snapshot pair that starts with a snapshot tag that corresponds to the most recently restored or applied snapshot to that file system. Exemplary systems identify the last applied or restored snapshot and request from the other file system the set of changes that corresponds to the changes made since the last applied or restored snapshot" (Paragraph 132). The examiner notes that Eshel teaches **"activating a migration procedure to copy source data files from the source storage device to locations of the corresponding target data files in the target storage device"** as "These embodiments of the present invention create a hot standby file system by first generating a snapshot of the original (source) file system and transferring the entire data set for that snapshot to a second file system in order to create an identical copy of the original file system (i.e., a mirror file system). These embodiments then periodically bring the standby or mirror file system up-to-date by generating new snapshots of the original file system and determining the changes between these new, more recently captured or generated snapshots and the state that was captured by a previous snapshot of the original file system that had been transferred to the mirror file system. The original file system generates a set of changes that are then communicated and applied to the standby file system in order to bring the standby file system up to the state of the new snapshots captured on the original file system" (Paragraph 130).

***Claim Rejections - 35 USC § 103***

12. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

13. Claims 22, 26, 55, and 59 are rejected under 35 U.S.C. 103(a) as being unpatentable over **Eshel et al.** (U.S. PGPUB 2003/0158862) as applied to claims 19-21, 23-25, 52-54, 56-58, 79, 88, and 90 above, and further in view of **Prahlad et al.** (U.S. PGPUB 2006/0010154).

14. Regarding claims 22 and 55, **Eshel** does not explicitly teach a method and system comprising:

A) wherein the target storage device comprises a NAS filer.

**Prahlad**, however, teaches “**wherein the target storage device comprises a NAS filer**” as “A NAS device may include a specialized file server or network attached storage system that connects to the network. A NAS device often contains a reduced capacity or minimized operating and file management system (e.g., a microkernel) and normally processes only input/output (I/O) requests by supporting common file sharing protocols such as the Unix network file system (NFS), DOS/Windows, and server message block/common Internet file system (SMB/CIFS). Using traditional local area network protocols such as Ethernet and transmission control protocol/internet protocol (TCP/IP), a NAS device typically enables additional storage to be quickly added by connecting to a network hub or switch” (Paragraph 12) and “The present invention provides, among other things, systems and methods for performing storage operations for electronic data in a computer network on a network attached storage device (NAS). Some of the steps involved in one aspect of the invention may include receiving electronic data from a network device for writing to the NAS device; writing the electronic data to the NAS device in a first location (i.e., primary storage); subsequently storing the electronic data to a second location (i.e., secondary storage); and storing a stub file at the first location, the stub file including a pointer to the second location that may redirect the network device to the second location if an access request for the electronic data is received from the network device” (Paragraph 17).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teachings of the cited references because teaching **Prahlad's** would have allowed **Eshel's** to provide a method to intercept file system calls in NAS devices, as noted by **Prahlad** (Paragraph 15).

Regarding claims 26 and 59, **Eshel** does not explicitly teach a method and system comprising:

A) wherein the selected information in a respective target data file identifies a physical location of the corresponding source data file on the source storage device.

**Prahlad**, however, teaches “wherein the selected information in a respective target data file identifies a physical location of the corresponding source data file on the source storage device” as “A stub file may contain some basic information to identify the file itself and also include information indicating the location of the data on the secondary storage device” (Paragraph 14).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teachings of the cited references because teaching **Prahlad's** would have allowed **Eshel's** to provide a method to intercept file system calls in NAS devices, as noted by **Prahlad** (Paragraph 15).

15. Claim 89 is rejected under 35 U.S.C. 103(a) as being unpatentable over **Eshel et al.** (U.S. PGPUB 2003/0158862) as applied to claims 19-21, 23-25, 52-54, 56-58, 79, 88, and 90 above, and further in view of **George** (U.S. Patent 6,993,679).

16. Regarding claim 89, **Eshel** does not explicitly teach a method comprising:

A) pausing the migration procedure after the request is received, based at least in part on an availability of resources; and

B) retrieving the requested data during the pause from a selected data file.

**George**, however, teaches “**pausing the migration procedure after the request is received, based at least in part on an availability of resources**” as “Note that in an alternative embodiment of a method of performing a backup, the backup procedure may progress normally, without first checking for the existence of any entries in the non-read list. If a read error is detected and the read error indicates that the address of the attempted read is on the non-read list, the backup may be paused and the data at that address may be restored (e.g., a system administrator may restore the data from another backup disk). Once the data has been restored and the address has been cleared from the non-read list, the backup may proceed” (Column 7, lines 14-23),

and “retrieving the requested data during the pause from a selected data file” as “Note that in an alternative embodiment of a method of performing a backup, the backup procedure may progress normally, without first checking for the existence of any entries in the non-read list. If a read error is detected and the read error indicates that the address of the attempted read is on the non-read list, the backup may be paused and the data at that address may be restored (e.g., a system administrator may restore the data from another backup disk). Once the data has been restored and the address has been cleared from the non-read list, the backup may proceed” (Column 7 , lines 14-23).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teachings of the cited references because teaching **George's** would have allowed **Eshel's** to provide a method to prevent requested data from being corrupt in a migration procedure, as noted by **George** (Column 2 , lines 18-22).

### ***Response to Arguments***

17. Applicant's arguments with respect to claims 19-26, 52-59, 79, and 88-90 have been considered but are moot in view of the new ground(s) of rejection.

### ***Conclusion***

11. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

Article entitled "Data Migration Solution" by **Falconstor** on 23 January 2003. The subject matter disclosed therein is pertinent to that of claims 19-26, 52-59, 79, and 88-90 (e.g., methods for data migration using stub files with NAS devices).

U.S. Patent 5,564,037 issued to **Lam** on 08 October 1996. The subject matter disclosed therein is pertinent to that of claims 19-26, 52-59, 79, and 88-90 (e.g., methods for data migration using stub files with NAS devices).

U.S. Patent 5,991,753 issued to **Wilde** on 23 November 1999. The subject matter disclosed therein is pertinent to that of claims 19-26, 52-59, 79, and 88-90 (e.g., methods for data migration using stub files with NAS devices).



U.S. PGPUB 2005/0015409 issued to **Cheng et al.** on 20 January 2005. The subject matter disclosed therein is pertinent to that of claims 19-26, 52-59, 79, ands 88-90 (e.g., methods for data migration using stub files with NAS devices).

U.S. Patent 7,103,740 issued to **Colgrove et al.** on 05 September 2006. The subject matter disclosed therein is pertinent to that of claims 19-26, 52-59, 79, ands 88-90 (e.g., methods for data migration using stub files with NAS devices).

U.S. PGPUB 2005/0033800 issued to **Kavuri et al.** on 10 February 2005. The subject matter disclosed therein is pertinent to that of claims 19-26, 52-59, 79, ands 88-90 (e.g., methods for data migration using stub files with NAS devices).

U.S. Patent 7,263,590 issued to **Todd et al.** on 10 February 2005. The subject matter disclosed therein is pertinent to that of claim 89 (e.g., methods to pause data backups).

#### ***Contact Information***

13. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Mahesh Dwivedi whose telephone number is (571) 272-2731. The examiner can normally be reached on Monday to Friday 8:20 am – 4:40 pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Tim Vo can be reached (571) 272-3642. The fax number for the organization where this application or proceeding is assigned is (571) 273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/Tim T. Vo/

Supervisory Patent Examiner, Art Unit 2168

Mahesh Dwivedi  
Patent Examiner

Art Unit: 2168

Art Unit 2168

March 5, 2008

/Mahesh H Dwivedi/

Examiner, Art Unit 2168